

REMARKS

Claims 1-24 are pending and are unamended. Applicant has amended the specification and the drawings to correct informalities. For the reasons stated below, Applicant respectfully submits that all claims pending in this application are in condition for allowance.

Specification and Drawing Objections

FIG. 1 was objected to because it includes reference numbers not mentioned in the description. In response, the highlighted reference numbers were deleted.

The specification was objected to for two informalities. In response, the specification was amended to correct the noted items.

Applicant believes that the amendments made to the specification and FIG. 1 overcome the objections stated in the Office Action. Applicant, therefore, respectfully requests that the Examiner withdraws the objections made to the specification and the drawings.

Prior Art Rejections

Claims 1, 2, 9, 12, 13, 20, 21, and 24 were rejected under 35 U.S.C. 102(e) as being anticipated by Eng et al. (hereafter, "Eng"), and claims 3-8, 10-11, 14-19, and 22-23 were rejected under 35 U.S.C. 103(a) as being unpatentable over Eng in view of Holden. These grounds of rejection are respectfully traversed.

1. Eng

Generally speaking, the purpose of Eng's invention is to provide a packet switch that can be modularly grown as large as an expanding network may require. As shown in Fig. 2 of Eng, a packet switching arrangement is formed by interconnecting M designated small-sized J×K output packet switches to separate outputs of an N×L interconnect fabric, where $N < L$ and $M = L/J$. The interconnect fabric further comprises a first stage of M P×J input modules, where $M = N/P$ and

P<J, and a second stage of J M×M intermediate modules. The interconnect fabric implements a routing algorithm for conveying the packets from the input modules to the output packet switches through the use of the intermediate modules.

In operation of the routing algorithm, packets arrive simultaneously in a given time slot at the inputs to the input modules. The time slot is mainly divided into M mini-slots, where M is the number of input modules. For each mini-slot, a separate particular output packet switch is assigned, or mapped, to each input module. Each mini-slot uses a distinct mapping within the same time slot. Each input module then examines the destination address in each packet received during the time slot and determines which packets are destined for the particular output packet switch assigned to the input module during that particular mini-slot. (Column 4, line 9 to column 7, line 58)

Eng differs from Applicant's invention in both its structure and routing mechanism.

a. Structure:

As shown in Fig. 2 of Eng, the structure of Eng's invention is essentially a three-stage interconnection network, where each stage is composed of a plurality of smaller switching modules. They are interconnected in such a way that each switching module in the first stage is connected to each of the switching modules in the second stage through a single path, and then similarly each switching module in the second stage is connected to each of the switching modules in the third stage through a single path. On the other hand, as shown in Fig. 1 of Applicant's invention, the structure of Applicant's invention is composed of an input module, an intermediate packet buffer, and an output module, where the input module and output module are respectively a single switch such as a crossbar switch, while the intermediate packet buffer is a single large block of memory composed of registers. Therefore, Applicant's structure is not an interconnection network such as disclosed in Eng.

b. Routing:

Firstly, Eng's routing mechanism uses "time-multiplexed switching" (referring to the division of each frame time into a number of mini-slots, one for each input switching module) and "route hunting" (referring to the seeking of an available intermediate switching module for

connecting each individual pair of input and output switching modules during each mini-slot (Column 7, line 59 to column 8, line 44)), both of which are considered too complicated for today's high-speed switching. In contrast, Applicant's routing mechanism does not involve such complications.

Secondly, according to Eng (Abstract, lines 9-18), only up to J simultaneous packets can be routed to separate inputs of a particular output packet switch for distribution to their respective destinations, while all other simultaneous packets destined for user equipment associated with the same output packet switch are lost. Moreover, according to Eng (Column 8, lines 12-21), packet loss will also happen during the "route hunting" process. That is, in each mini-slot, for any given pair of input and output switching modules, if no available intermediate switching module can be found to simultaneously connect the two, packets between the two will be lost. On the other hand, in Applicant's invention, packet loss will occur only if the intermediate packet buffer is all full. Therefore, the probability of packet loss is much smaller in Applicant's invention compared to the scheme in Eng.

Therefore, the two inventions of Eng and Applicant are fundamentally different.

2. Holden

Holden's invention relates to an ATM switching system architecture whose structure is an interconnection network of smaller ATM switch elements (see Fig. 2 of Holden), whereas Applicant's invention relates to a general packet switching system whose structure is an improved shared-buffer memory switch (see Fig. 1 of Applicant's disclosure). Although Holden also has the ingredient of shared-buffer memory, shared-buffer memory by itself is well-known. Holden's invention has no improvement on the shared-buffer memory architecture but only focuses on the control mechanism including techniques such as "linked list RAM and multipriority queues" (Columns 6-7), "back pressure" (Columns 7-8), "aggregate bits" (Column 9), "proportional bandwidth queues" (Column 10), "multicasting" (Columns 11-12), and "marked interrupt linked list" (Column 13). As a result, since Holden still uses conventional shared-buffer memory, Holden's switching architecture also has the same problems faced by any other architectures employing conventional shared-buffer memory. As discussed in Applicant's

specification, the throughput of a conventional shared-buffer memory switch is limited by the bandwidth of the memory bus, which is also potentially a single point of failure. Moreover, the cost of memory-bus operations like multiplexing, demultiplexing, serial-to-parallel (S/P) conversion, and parallel-to-serial (P/S) conversion is high. Therefore, the shared-buffer memory style of switching architecture does not scale well for large number of broadband I/O ports. This explains why Holden only uses a very small "shared-buffer memory" (Holden's alternate terminology for "shared-buffer memory" in his specification is "cell buffer pool memory") in his design (Column 5, lines 59-60). Holden's cell buffer pool memory has a typical size of 32 individual cell memories (Column 6, line 19) supporting only 8 input lines and 8 output lines, where each input/output line is a 4-bit wide memory bus (Column 6, lines 6-9). Moreover, multiplexers are used for connecting input buses to the cell memories according to Holden's architecture (Column 6, lines 24-25).

On the other hand, Applicant's invention improves the conventional shared-buffer memory architecture in such ways that (1) the access to the memory storage in which the packet buffer resides is not through a memory bus or buses, since a memory bus would only allow the access of addresses one at a time, while the memory storage adopted in Applicant's invention allows massive parallel access; and (2) by employing circular shift registers in the buffer memory, the access to the memory storage also does not require multiplexing and demultiplexing, and the packets can be pipelined into and out of the memory switch in serial bits throughout their journey without the need of serial-to-parallel and parallel-to-serial conversion of the packet format, thereby greatly simplifying the operation.

3. Patentability of claims 1, 12, 21, and 24 over Eng:

a. Claim 1

Claim 1 reads as follows:

An M×N packet switch for switching M input packets arriving in each of a sequence of frame times to N output ports, the switch comprising

an input module, having M inputs and B outputs, $B > M$, for switching the M input packets to M of the B outputs to produce M switched packets during each of the frame times,

a packet buffer including B registers, coupled to the input module, for storing the M switched packets into M available registers during each of the frame times to produce M stored packets, and

an output module, having B inputs and N outputs coupled to the packet buffer, for transferring up to N packets from occupied registers in each of the frame times to the output ports based upon destination addresses contained within each of the stored packets. (underlining added for emphasis)

The Examiner asserts that Eng discloses the subject matter of Applicant's claim 1, particularly at the Abstract, lines 2-9; column 1, lines 48-59; column 2, lines 60-65; column 6, lines 41-46; and column 6, line 60 through column 7, line 3. This is incorrect. In fact, as described above, the system disclosed by Eng is an interconnection network, which is a completely different system than Applicant's claimed invention that is an improved shared-buffer memory switch. In particular, the following portions of Eng were highlighted by the Examiner:

(i) Eng's Abstract, lines 2-9 reads as follows:

The invention comprises an N input, L output interconnect fabric ($L > N$), and a plurality of $J \times K$ smaller packet switches ($J > K$). Each of the J inputs to each packet switch is connected to a separate one of the L outputs of the interconnect fabric, and each of the K outputs from each packet switch is connected to a destination equipment. (underlining added for emphasis)

The underlined parts clearly state that the overall structure of their inventive system is an interconnection network.

(ii) Eng's column 1, lines 48-59 merely repeats the disclosure of the structure of the interconnection network system.

(iii) Eng's column 2, lines 60-65 discloses one inventive aspect of the routing process employed in their system, that is, mapping the address in each packet to a subset of several interconnect fabric outputs, and then sending the packet to any available output in that subset. This is not relevant to the subject matter in Applicant's claim 1. Moreover, as discussed above, the routing process employed by Eng and that employed by Applicant are fundamentally different.

(iv) Eng's column 6, lines 41-46 describes the particular details of each of the input modules of the interconnect fabric of their inventive system. There is nothing whatsoever in this description that discloses or suggests the subject matter in Applicant's claim 1. Although Applicant's claim 1 contains the term "input module" as well, the term by itself is generic, and, by referring to the two specifications respectively, the Applicant's "input module" and Eng's "input module" are referring to different subject matter. Referring to Figs. 2 and 3 of Eng, the system in Eng comprises a number of input modules, where each input module is a small node of the overall interconnection network. On the other hand, referring to Fig. 1 of Applicant's invention, Applicant's system contains only one input module, which is part of the claimed packet switch. Therefore, Eng's column 6, lines 41-46 is also not relevant to Applicant's claimed subject matters.

(v) Eng's column 6, line 60 through column 7, line 3 describes the first step of the routing process in each mini-slot. As mentioned before, the routing processes respectively employed by Applicant and Eng are fundamentally different.

In sum, a switch comprising the combination of highlighted elements as recited in claim 1 is not disclosed or suggested in Eng, and therefore claim 1 is deemed allowable over Eng.

b. Claims 12, 21 and 24

Claims 12, 21 and 24 are patentable for at least the same reasons as given above with respect to claim 1.

Claim 12 recites, in part:

input means, having M inputs and B outputs, $B > M$, for switching the M input packets to M of the B outputs to produce M switched packets during each of the frame times,

storage means, including B registers, coupled to the input module, for storing the M switched packets into M available registers during each of the frame times to produce M stored packets, and

output means, having B inputs and N outputs coupled to the packet buffer, for transferring up to N packets from occupied registers in each of the frame times to the output ports based upon destination addresses contained within each of the stored packets
(underlining added for emphasis)

Each of claims 21 and 24 recites, in part:

switching the M input packets to M of the B outputs to produce M switched packets during each of the frame times, $B > M$,

storing the M switched packets into M of B registers during each of the frame times to produce M stored packets, and

transferring up to N packets from up to N of the B registers in each of the frame times to the output ports based upon destination information
(underlining added for emphasis)

Again, by way of reiteration, Eng's disclosures, in particular, Abstract, lines 2-9; column 1, lines 48-59; column 2, lines 60-65; column 6, lines 41-46; and column 6, line 60 through column 7, line 3, do not teach or suggest the subject matters, especially the underlined parts, claimed in Applicant's claims 12, 21 and 24.

4. Patentability of claims 10 and 23 over Eng in view of Holden

a. Claim 10

Claim 10 recites, in part:

[(a)] an M×B input crossbar switch, $B > M$, for switching the M input packets to M of the B outputs to produce M switched packets during each of the frame times,

[(b)] a one-stop shared buffer memory, including B registers, coupled to the input crossbar switch, for storing the M switched packets into M available registers during each of the frame times to produce M stored packets,

[(c)] a B×N output crossbar switch coupled to the packet buffer, for transferring up to N packets from occupied registers in each of the frame times to the output ports based upon destination addresses,

[(d)] a register selector for assigning the M of the B registers during each of the frame times to generate M assigned registers,

[(e)] M header hoppers, coupled to the input crossbar switch, for storing header information from each of the M input packets in each of the frame times and M addresses of the M assigned registers for the M input packets in each of the frame times
(underlining added for emphasis, elements (a)-(e) added for reference purposes)

The Examiner asserts that Eng discloses the subject matter of the parts (a) and (c) of Applicant's claim 10, particularly at the Abstract, lines 2-9; column 1, lines 48-59; column 2, lines 60-65; and column 6, line 60 through column 7, line 3. Referring to the detailed discussion with respect to claim 1, Eng does not teach or suggest anything related to the subject matter of parts (a) and (c) of Applicant's claim 10.

The Examiner further asserts that Holden discloses the subject matter of the parts (b), (d) and (e) of Applicant's claim 10, where, in particular,

(1) Holden's Abstract, lines 4-6 and column 6, lines 41-44 discloses part (b),

(2) Holden's column 5, lines 58-62 and column 15, lines 24-26 discloses part (d), and
(3) Holden's column 5, lines 39-44 and 58-62 and column 16, lines 32-40 discloses part (e).
This is incorrect. In fact, as mentioned above, Applicant's invention focuses on improvements over conventional shared-buffer memory switch architecture, whereas Holden has no improvement on conventional shared-buffer memory and employs conventional shared-buffer memory as one of his components.

(1) Holden's abstract, lines 4-6 reads as follows:

A shared pool of memory is employed to eliminate the need to provide memory at every crosspoint.

The claim element in Applicant's claim 10, part (b) is the "one-stop shared buffer memory", which is different from Holden's "shared pool of memory". According to Applicant's specification, a "one-stop shared buffer memory", in contrast with conventional shared-buffer memory like the one used by Holden, means a packet buffer such that whenever a packet occupies a certain section, e.g., a register, in the buffer, it remains in that section until its eventual exit from the switch. In fact, no one-stop shared buffer memory is even mentioned in either Holden or Eng. Holden's column 6, lines 41-44 describes the use of linked list RAM to maintain its FIFO queues. Linked list RAM is also not present in, and completely irrelevant to, Applicant's claimed subject matter.

(2) Holden's column 5, lines 58-62 emphasizes the use of a very small shared-buffer memory in each of its switch element circuits, while Holden's column 15, lines 24-26 is a part of his claim 6, including a step of "receiving cells and storing those cells in a cell memory". Again, this disclosure is not relevant to the claimed "register selector" (the underlined claim element in Applicant's claim 10, part (d)). One preferred embodiment of a register selector is described on page 10, line 19 through page 11, line 13 of Applicant's specification.

(3) Holden's column 5, lines 39-44 describes the usage of a connection table controller for reading header information from the workstation interface and using that header information to add an appropriate switch tag to the cells before they are transmitted to the switch fabric. Applicant's claimed subject matter does not include such controllers. Holden's column 5, lines 58-62 has been discussed in the preceding text. Holden's column 16, lines 32-40 is again a part of his claim reciting a number of steps including "receiving and storing cells," "using the controller to examine a tag portion of a cell," and "using the controller to queue the cell by adding an identifier." None of this discussion in Holden is related to the claimed "header hopper" (the underlined claim element in Applicant's claim 10, part (e)). Neither Eng nor Holden has a component equivalent in function to Applicant's "header hopper." One preferred embodiment of a header hopper is described on page 8, line 19 through page 9, line 10.

In sum, Applicant's claim 10 is not disclosed or suggested by the combination of Eng and Holden. Therefore, claim 10 is deemed allowable over this combination of references.

b. Claim 23

Claim 23 recites, in part:

[(e)] sending headers from the M input packets to the header hoppers,

[(g)] updating the queues based on the header information provided by the header hoppers,

[(h)] sending control information to a register selector to inform the register selector of the destination addresses of the M input packets in each of the frame times,

[(j)] transmitting the up to N selected stored packets to the outputs,
updating the register selector to account for any remaining destination
addresses for each stored packet
(underlining added for emphasis, elements (e)-(j) added for reference purposes)

Application No. 09/882,760
Reply to Office Action of January 11, 2005

The Examiner alleges that Eng's Abstract, lines 2-9;, column 1, lines 48-59;, column 2, lines 60-65; and column 6, line 60 through column 7, line 3, together with Holden's Abstract, lines 4-6;, column 5, lines 39-44 and 58-62; column 6, lines 41-44, column 15, lines 24-26; and column 16, lines 32-40 discloses the claimed subject matter of Applicant's claim 23. This is also incorrect. In view of the above discussion, at least steps (e), (g), (h), and (j) are not disclosed or suggested by Eng or by Holden. Accordingly, claim 23 is believed to be patentable over Eng in view of Holden.

5. Patentability of dependent claims:

The dependent claims are patentable for at least the reason that they are dependent upon allowable independent claims and because they recite additional patentable features.

Conclusion

Insofar as the Examiner's rejection has been fully addressed, the instant application is in condition for allowance. Issuance of a Notice of Allowability of all pending claims is therefore earnestly solicited.

Respectfully submitted,

SHUO-YEN ROBERT LI et al.

May 11, 2005
(Date)

By:

Clark Jablon

CLARK A. JABLON

Registration No. 35,039

AKIN GUMP STRAUSS HAUER & FELD LLP

One Commerce Square

2005 Market Street - Suite 2200

Philadelphia, PA 19103

Direct Dial: (215) 965-1293

Facsimile: (215) 965-1210

Application No. 09/882,760
Reply to Office Action of January 11, 2005

Amendments to the Drawings:

The attached formal drawing sheet replaces the original sheet of formal drawings.

Fig. 1: Delete label "116" from the lead line (path) between elements 112 and 132.

Delete label "141" from the lead line (path) between elements 106 and 134.

Attachment: Replacement Drawing Sheet (1 sheet)